

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 124 173 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
16.08.2001 Bulletin 2001/33

(51) Int Cl.7: G06F 1/00, H04L 29/06

(21) Application number: 01201844.6

(22) Date of filing: 18.05.1998

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE

(30) Priority: 20.05.1997 US 47235 P

(62) Document number(s) of the earlier application(s) in
accordance with Art. 76 EPC:
98922326.8 / 0 983 540

(71) Applicant: America Online, Inc.
Dulles, VA 20166 (US)

(72) Inventors:
• Morris, Harry W.
Reston, Virginia 20190 (US)

• Lippke, David Lowell
Herndon, Virginia 20170 (US)
• Bosco, Eric
Arlington, Virginia 22201 (US)
• Steele, Colin Anthony
Leesburg, Virginia 20175 (US)

(74) Representative: Charig, Raymond Julian
Eric Potter Clarkson,
Park View House,
58 The Ropewalk
Nottingham NG1 5DD (GB)

Remarks:

This application was filed on 16 - 05 - 2001 as a
divisional application to the application mentioned
under INID code 62.

(54) Self-policing, rate limiting online forums

(57) A method and means for self-policing and automatically rate-limiting multiple-user online forums. The preferred embodiment of the invention includes a set of rules that permit users to censure other users. A censured user has one or more "privileges" (which may include access to the online computer system or the effective message rate of a communications connection) taken away or diminished; the privilege is gradually restored if the censure user behaves. The censoring participant and other pertinent participants are notified of the effect of the censoring on the censured user. In another aspect of the invention, the online computer system automatically tracks the rate at which a user sends certain types of messages, and can message rate limit a user who uses too many system resources by sending a large number of messages in rapid succession.

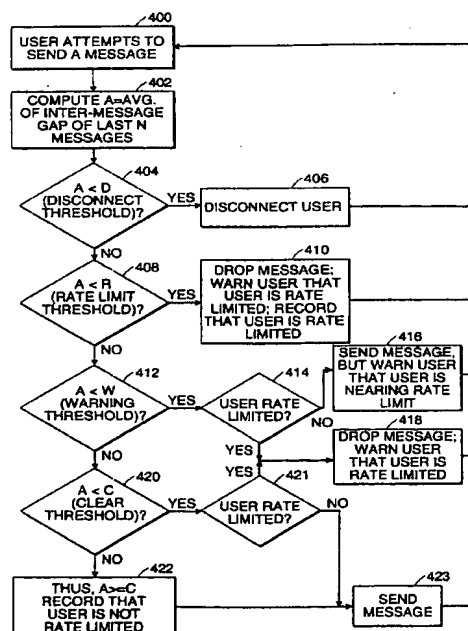


FIG. 4

EP 1 124 173 A2

Description**Technical Field**

5 [0001] This invention relates to distributed computer services, particularly computer services having online forums.

Background Information

10 [0002] An online forum is a communications interchange in which people may communicate with others through successive electronic transmissions between respective computer systems. An online forum, or any other type of distributed computer services, may be implemented on a distributed computer system such as that shown in FIG. 1. Forum participants (equivalently, users of the computer services) typically are scattered across a large geographical area and communicate with one or more central server systems 100 through respective client systems 102 (e.g., a personal or laptop computer). In practice, the server system 100 typically will not be a single monolithic entity but rather
 15 will be a network of interconnected server computers, possibly physically dispersed from each other, each dedicated to its own set of duties and/or to a particular geographical region. In such a case, the individual servers are interconnected by a network of communication links, in known fashion. One such server system is "America Online" from America Online Incorporated of Virginia.

[0003] Each client system 102 runs client software that allows it to communicate in a meaningful manner with corresponding software running on the server system 100.

[0004] The client systems 102 communicate with the server system 100 through various channels, such as a modem 104 connected to a telephone line 106 or a direct Internet connection using a transfer protocol such as TCP/IP (Transfer Control Protocol/Internet Protocol). The server system 100 is responsible for receiving input from the client systems 102, manipulating the collective body of input information (and possibly information from other sources) into a useful
 25 format, and retransmitting the formatted information back to one or more clients 102 for output on an output device, such as a display screen.

[0005] Referring to FIG. 2, one type of forum is a "chat room" 200, in which the various participants 204 (e.g., "Allens9," "JOSHUA ALEX," etc.) may enter text which appears in a scrolling text window 202 on each participant's computer display screen. In the example in FIG. 2, the chat room 200 has 22 participants whose identities (or "screen names") are listed in a scrolling window 210. A participant 204 may respond to the comment of another participant 204 by
 30 entering a line of text in an edit box 206 and activating (e.g., by clicking with a pointer device, such as a mouse) a SEND button 208. In response, the text in the scrolling text window 202 scrolls upwards and the newly entered line of text is displayed at the bottom of the scrolling text window 202. In the illustrated example, the last participant to enter a comment was JOSHUA ALEX, who typed "TEXAS."

35 [0006] The chat room 200 shown in FIG. 2 is "public", meaning that it has multiple participants who were placed in the chat room by the computer-service provider and who most likely never have met or conversed with one another before. A comment by a participant in a public forum may be seen by all of the participants of the chat room. If a participant desires some privacy, that participant may "open" and enter a "private" chat room (for example, by clicking on a SETUP button 212), and thereafter invite one or more other participants to enter the private chat room. Once in
 40 a private forum, participants may communicate with one another without fear that uninvited participants will be able to see their comments.

[0007] When a participant in a forum, whether public or private, makes a comment that others in the forum regard as offensive, in poor taste, wildly incorrect, or otherwise objectionable, the offending participant most likely will be "flamed" by one or more of the other participants. A "flame" is a reprimand or other stringent response directed at the
 45 offending party. One purpose behind flaming another participant is to dissuade the offender, through embarrassment or intimidation, from making further objectionable comments. In this manner, if the offending user chooses to curb his or her behavior in response to the flaming, a forum may be crudely regulated or "policed" by the forum's participants. However, the offending participant may continue to behave in an objectionable manner. Further, a participant who overly "flames" other participants may also be objectionable. Accordingly, participant policing of forums does not always
 50 work well. In such cases, offended participants may drop out of "flame-filled" forums, and/or the online service must devote resources to actively police problematic participants.

[0008] Other objectionable behavior includes sending one or more messages to "spoof" other users as to the sender's identity in order to try to get confidential information (e.g., credit card numbers or passwords) sent in response (sometimes called "password fishing").

55 [0009] Another problem that can arise in online systems is "resource hogging", where a participant uses features such as broadcast or multi-cast messaging to send a large number of messages to other users in a short period of time (sometimes called "spamming"). Such resource hogging deprives other users of server resources, and can slow an online system response time to undesirable levels.

[0010] Accordingly, the inventor has determined that there is a need for a better way to police recalcitrant participants in online forums and to reduce spamming. The present invention provides a method and means for accomplishing this goal.

5 SUMMARY

[0011] The invention provides a method and means for self-policing and automatically rate-limiting multiple-user online forums. The preferred embodiment of the invention includes a set of rules that permit users to censure other users. A censured user has one or more "privileges" (which may include access to the online computer system or the effective message rate of a communications connection) taken away or diminished; the privilege is gradually restored if the censured user behaves. The censuring participant and other pertinent participants are notified of the effect of the censuring on the censured user. In another aspect of the invention, the online computer system automatically tracks the rate at which a user sends certain types of messages, and can message rate limit a user who uses too many system resources by sending a large number of messages in rapid succession. The amount of rate limiting may be a function of the amount of censure that has been applied to the user being limited.

[0012] Advantages of this invention may include one or more of the following. The techniques described here enable a multiple-user online service (e.g., a chat room or other forum) to be self-policing. Access to the computer service for a particular user is automatically regulated by the computer based on input from other users concerning the conduct of the user under consideration, or based on the message rate of a user. Users of a computer-based system have the ability to sanction a misbehaving user and thereby have the offending user's access to the system denied or curtailed. Unlike the conventional "flaming" approach to policing, which typically fails because it relies on the misbehaving user curbing his or her own behavior or results in retaliatory "flaming", the policing techniques described here are based on predetermined rules and operate automatically in response to votes cast by other users of the computer system. The offending user has no choice in the matter once he or she chooses to misbehave. Because the operation of these policing techniques are automatic, the computer-based service provider need not expend personnel time and resources to police public forums, direct user communications, and the like.

[0013] The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

30 DESCRIPTION OF DRAWINGS

[0014] FIG. 1 shows a prior art distributed computer system of the type used for providing online computer services.

[0015] FIG. 2 is a screen shot showing an example of a prior art online computer forum.

[0016] FIG. 3 is a flowchart of a basic embodiment of the self-policing aspect of the invention.

[0017] FIG. 4 is a flowchart of a basic embodiment of the rate-limiting aspect of the invention.

[0018] Like reference numbers and designations in the various drawings indicate like elements.

40 DETAILED DESCRIPTION

[0019] The prospect of millions of concurrent users of an online computer system puts great pressure on the ability of the system provider to police abusive participants. Accordingly, the invention provides a self-policing online environment that makes participants responsible for their behavior. That is, other participants can express an opinion about that behavior, and have that opinion affect the offending user in a negative way and be "visible" to other participants. Providing such a self-policing user feedback mechanism lowers the cost of providing online services to users and provides a more "user-friendly" environment for participants.

[0020] In another aspect of the invention, the online computer system automatically tracks the rate at which a user sends certain types of messages, and can "rate limit" a user who "hogs" too many system resources by sending a large number of messages (or messages of selected types) in rapid succession.

Participant Self-Policing

[0021] FIG. 3 is a flowchart of a basic embodiment of the self-policing aspect of the invention. Initially, user A generates an event I, such as a message (STEP 300). Event I is transmitted to some number of other participants, such as users B and C in a public forum on an online computer system (STEP 302). However, event I may be a direct communication between user A to user B, such as by use of the Instant Message™ feature of the America Online computer service. User B may feel that event I is "objectionable" or "evil" (a subjective decision by B), and therefore generates a "vote" against the behavior of user A by sending a special type of response message - an "evil" event E

(STEP 304). In the preferred embodiment, a user cannot "evil" another participant except in response to a message from the participant, or otherwise in cases where the participant's actions affect the user doing the "eviling." In order to reduce retaliation, users cannot "evil" other participants directly in response to being "eviled".

[0022] The online system receives event E and checks a database to see whether user B has the right to "evil" user A (STEP 306). For example, this step can limit users who themselves have been restricted from "eviling" other users.

[0023] If user B does not have the right to "evil" user A, then user B is notified of the effect of user B's action on user A's "evil index" - a value of how restricted user A is with respect to use of resources on the online system (STEP 308). In this case, user B's action has no effect on user A.

[0024] If user B does have the right to "evil" user A, then user A's "evil index" is modified by a suitable amount, which will affect further actions by user A (STEP 310), as described below. The amount of modification can be based on the past behavior of user A, or of users A and B. The amount of modification may also depend on the type of "eviling" asserted by user B. For example, user B may wish to assert an "eviling" event E anonymously rather than be identified. Anonymous "eviling" may be accorded lesser weight. That is, an anonymous "eviling" response may count as fewer evil "votes" than it would if the eviling user's identity was revealed. In one embodiment, an "eviling" user may set up an automatic "eviling" response to all communications from another specific user. In this case, the effect of the "eviling", which is intended to discourage communication from the specific user, is given very low weight (equivalently, counts as very few evil "votes") because of its automatic invocation.

[0025] User A's "evil index" is "broadcast" to other users who "know about" user A (STEP 312). For example, such users might be all of the participants in a private chat room, or all visitors to a public chat room. In one embodiment, a modified "evil index" is reported to each user in a current chat room's user list and to each user who has the "eviled" users on his or her "buddy list". Thus, an "eviled" user is publicly chastised.

[0026] Thereafter, user A's "evil index" begins to gradually "decay" back to normal over time (STEP 314). Such decay may be linear, exponential, step-wise, or some other function. Meanwhile, user B is notified of the effect of user B's action on user A's "evil index" (STEP 308). In this case, user B's action has had an effect on user A.

[0027] Some of the steps noted above may be done in different order without substantially changing the effect of the process. For example, STEPS 312, and 314, and 308 may be performed in different order.

[0028] A basic "penalty" for having a non-normal "evil index" is denial of access to a forum or the online service until the user's "evil index" has decayed back to normal. In a more refined embodiment, a user's "evil index" affects a rate limit which governs a user's ability to send (and/or receive) messages. This feature allows other participants to "evil" a user who "flames" or "spams" them, and thus reduce the rate at which the recalcitrant user can send and/or receive messages. A description of rate limiting is set forth below.

[0029] A server database (which may be centralized or distributed) stores a user's "evil index". A user's "evil index" can be maintained in a user-record as a global total, or by forum, or both. The value of each "evil index" can be used to control the user's ability to log on to the online computer system or access selected forums, and/or the effective rate of message or file transmissions.

[0030] Other rules may applied in different embodiments. For example, the following rules can be implemented:

A user must be in a forum (e.g., a chat room, but including direct communication between users, such as the Instant Message™ feature of the America Online computer service) for a specified amount of time before being allowed to "evil" another user in that forum; this reduces "hit-and-run flaming". The minimum amount of time may vary from forum to forum, and from user to user (e.g., a forum "sysop" may be immune to "eviling"). A user's user-record in the server database thus would record a time-of-entry for a forum. For example, a user's time-of-entry to a forum would be compared to the current time in STEP 306 of FIG. 3 to determine if the user had the right to "evil" another participant.

A user must be in a forum for a specified amount of time before being allowed to "evil" another user who has been in that forum for a longer period of time. For example, a user's time-of-entry to a forum would be compared to the time-of-entry of another participant in STEP 306 of FIG. 3 to determine if the user had the right to "evil" that other participant. The specified amount of time may vary from forum to forum, and from user to user.

A user's eviling response may be accorded increased weight (equivalently, counted as extra "evil" votes) based on the "seniority" of the "eviling" user. Each additional unit of time spent in a forum could enhance a user's "seniority", thereby allowing long term user's more "eviling" power than new-comers. That is, being "eviled" by an "old-timer" can have a greater than normal affect on modification of a chastised user's "evil index" in STEP 310 of FIG. 3. A user's user-record in the server database thus would record a total-time-of-access for each forum, which could be a "lifetime" total (such that users who had belonged to the online computer service for longer periods would have greater "eviling" power than relatively new subscribers to the service), a running total for a selected period of time, or a session total. The time period for accumulating extra votes and the "eviling" value of extra votes may vary

from forum to forum, and from user to user.

A user may be limited in the number of "eviling" votes that can be cast in any one online session or time period (e.g., a day or week). A user's user-record in the server database thus would record the number of "eviling" votes cast globally or by forum. For example, a user's number of previous "eviling" votes cast could be examined in STEP 306 of FIG. 3 to determine if the user had the right to "evil" another participant.

The amount of modification of a user's "evil index" in STEP 310 of FIG. 3 after being "eviled" is preferably non-linear, so that the first few times a user is "eviled" has little effect (and possibly no effect until some threshold level of accumulated "eviling" votes from several participants are asserted against the user) on the user's "evil index" (everyone makes a few mistakes). However, for subsequent times that a user is "eviled", the user's "evil index" preferably is modified to a much greater degree, increasing the penalty for recidivism. The determination of how much to modify a user's "evil index" can be based on the user's global "evil index" (i.e., a "rap sheet" evaluation) or forum "evil index" (i.e., a "fresh start" approach). A user's user-record in the server database would thus record the number of times the user has been "eviled", either globally or by forum, for example. The recorded number might be a "lifetime" total or a running total for a selected period of time. The amount of "evil index" modification may also vary from forum to forum, and from user to user.

As an example, each user may be given an initial "evil index" of 100. A first "offense" may reduce that value to 95; a second "offense" reduces the value to 85; third and subsequent offenses reduce the current value by 15 units. Alternatively, each user is given an initial "evil index" of 0 and "eviling" increases that value. If a range of 100 is used, a user's "evil index" may be regarded as a "percentage of evil", up to 100% "evil".

The decay rate for returning a user's modified "evil index" to normal can vary from forum to forum. For example, the decay in chat rooms (where most flaming occurs) may be less than in other forums. Using the example from immediately above, the user's "evil index" may "decay" back to "normal" at the rate of 2 units per hour in a chat room, but at the rate of 5 units per hour in other forums.

A user may query a forum to determine the rate of "eviling" within the forum as a measure of how much members of the forum "evil" one another. A high rate may indicate that one or more members are misbehaving. The system keeps track of all users participating in a forum, so the current "eviling" rate is a time-weighted average of the number of "eviling" votes cast by the current participants. In an alternative embodiment, a user may query a forum to determine the accumulated value of the "evil index" of all current users. This measures how often the users have misbehaved in the past, and can serve as an estimate of the likelihood that the current users will misbehave in the future. The system keeps track of all users participating in a forum, so the total "evil index" is the sum of the pertinent forum "evil index" for each participant.

In some embodiments, a user who has been "eviled" has a lessened ability to "evil" other users (thus reducing retaliation). However, some online systems implement message types, such as broadcast or multi-cast messages or self-repeating messages, that are more frequently used than other message types to flame or spam other participants. In some embodiments, it may be desirable to allow all users (even those with a modified "evil index", and thus possibly restricted in usage rights) to "evil" the originator of such message types. Such a determination would be made in STEP 306 of FIG. 3 by examining the message type of event I sent by the originating user. The amount of modification of such an originator's "evil index" can be set to be greater than for flaming to discourage use of such message types for spamming.

[0031] When a user's "evil index" is modified, the user is notified that privileges, such as rate of messaging, have been limited. In one embodiment, a message is sent from the server that contains the user's current "evil index" to the "eviled" user, and optionally includes the current rate of decay back to normal. Such information allows a wide range of feedback to be presented to the user about his or her ability to interact with the online computer system. For example, a graphical "power meter" or "power bar graph" can be used to indicate the "evil index" of the user. For instance, a color-coded bar graph can be divided into (1) a green zone to represent that the user's "evil index" is normal; (2) a yellow zone to represent that the user's "evil index" has been modified slightly; (3) a red zone to represent that the user's "evil index" has been modified significantly; and (4) a black zone to represent that access or message privileges have been suspended for a time. However, other methods of informing a recalcitrant user of his or her "evil index" can be used.

[0032] In one embodiment, a user can "rehabilitate" his or her "evil index" by visiting advertisements displayed on the online computer system; "visits" can be determined by requiring the user to "click" on an advertisement. The user

is rewarded with more "power" by adjusting the value of the user's "evil index" more towards normal.

Automatic Rate Limiting

5 [0033] In one embodiment, both input and output messaging rates of a user are limited based on the behavior of the user and/or available system resources. Such rate limiting can stop malicious users and errant client computers from consuming more than a fair share of online system resources. However, preferably the rate limiting system is weighted to tolerate brief bursts of messaging activity while penalizing unacceptably large rates of messaging. Rate limiting can also limit aggregate input to a server to a level at which the system is reasonably loaded under normal conditions. Rate limiting can also be combined with "eviling" by automatically adjusting a users rate limit parameters based on their "evil index."

10 [0034] In one embodiment, input rate limiting - from user to server - is conducted on a per user connection basis, and within a user connection on a per-message type bases (that is, rate limits for different types of messages may be set to different values). In one embodiment, rate limiting for a user is achieved in accordance with the following algorithm:

15 (1) Define A as the running average of inter-message time gaps for the last N messages of selected types that the user has attempted to send; a system selected value I is used as the first value for A . Calculation of A can be done, for example, on a forum basis (accounting only for messages sent in the current forum), session basis (accounting for messages sent in the user's current online session), or message-count basis (accounting for the last N messages sent at any time in any forum).

20 (2) If A is below a warning threshold W (indicating that the user is approaching a point of sending messages too frequently), when the user attempts to send a message, send the user a warning message but transmit the user's message.

25 (3) If A is below a rate limit threshold R (indicating that the user is sending messages too frequently) when the user attempts to send a message, send the user a warning message and drop the user's message.

(4) Repeat the above steps until A rises above a clear threshold C (indicating that the user is not sending messages too frequently), at which time the rate limiting condition is considered cleared.

(5) If at any time A drops below a disconnect threshold D , disconnect the user.

30 [0035] FIG. 4 is a flowchart of a basic embodiment of the rate-limiting aspect of the invention, showing a slightly different order for the steps described above. A user attempts to send a message (STEP 400). Average A is computed (STEP 402). If A is less than a disconnect threshold D (STEP 404), the user is disconnected (STEP 406). Otherwise, if A is less than a rate limit threshold R (STEP 408), the message is dropped, the user is warned that the rate limited has been exceeded, and the user is flagged as "rate limited" (STEP 410).

35 [0036] Otherwise, if A is less than a warning threshold W (STEP 412), a determination is made as to whether the user is rate limited (STEP 414). If not, the message is sent, but the user is warned that the rate limit is being approached (STEP 416). Otherwise, the message is dropped and the user is warned that the rate limited has been exceeded (STEP 418).

40 [0037] If A is not less than the warning threshold W (STEP 412), and A is less than a clear threshold C (STEP 420), a determination is made as to whether the user is rate limited (STEP 421). If not, the message is sent (STEP 423). Otherwise, the message is dropped and the user is warned that the rate limited has been exceeded (STEP 418).

[0038] Finally, if A is not less than the clear threshold C (STEP 420), the rate limit flag for the user is cleared (STEP 422) and the message is sent (STEP 423).

45 [0039] The rate limiting algorithm supports several "tunable" parameters:

- The running average of inter-message time gaps - A
- The number of message receptions over which A is calculated - N
- An initial average - I
- A clear threshold - C
- 50 • A warning threshold - W
- A rate limit threshold - R
- A disconnect threshold - D

55 [0040] In one embodiment, the values for C , W , R , and D are selected such that $C > W > R > D$. The initial average rate I can be weighted to increase the algorithm's tolerance of bursts of activity, such as "chatty startups" when a user joins a forum and sends several messages in succession. The threshold rates can be set globally for all user's, or "tuned" for each user.

[0041] In one embodiment, the difference between the clear threshold C and the rate limit threshold R can be "tuned"

by the online computer system to alter the interval between commencement of rate limiting for a user and the resumption of normal activity. The difference between C and R may be automatically increased, for example, if the user sends excessive "spamming" type messages at too great a rate. Further, the threshold values C , W , R , and D can be dynamically modified by the online system as a way of limiting system resource usage if too many users are simultaneously using the system.

[0042] As noted above, when a user's message rate is limited, the user is notified. In one embodiment, a message is sent from the server to a rate limited user that contains values for each of the parameters outlined above, and another message is sent the next time that the server will accept messages from the user without restriction. Such information allows a wide range of feedback to be presented to the user about his or her ability to interact with the online computer system. For example, a graphical "power meter" or "power bar graph" can be used to indicate the "health" or "power" of the user. For instance, a color-coded bar graph can be divided into (1) a green zone to represent that no rate limiting is in effect; (2) a yellow zone to represent that the user's message rate is near the point where rate limiting would take effect; (3) a red zone to represent that message rate limiting is in effect; and (4) a black zone to represent that access privileges have been suspended for a time. However, other methods can be used to inform a recalcitrant user of his or her message rate limit status.

[0043] In one embodiment, a user can "rehabilitate" his or her message rate limit status by visiting advertisements displayed on the online computer system. The user is rewarded with more "power" by adjusting the parameters.

[0044] In one embodiment, output rate limiting - from server to user - is performed in a similar manner, with the following differences:

6) The system server audits the rate of outgoing error messages of selected types (e.g., RATE_TO_HOST), and uses the above rate limiting algorithm with $R=D$, to prevent excessive load due to handling users that are generating rate errors.

7) The system server can limit the rate of outgoing messages of selected types (e.g., those types that may be used for spamming) using the above rate limiting algorithm when $A < R$. In this case, a notice message can be sent to the user, to the user's intended recipient, and to a server storing the user's "evil index". The intended recipient can "evil" the sending user, or the system server can automatically modify the user's "evil index" based on the sending of excessive "spamming" type messages at too great a rate.

[0045] The algorithms described above prevent abuse of online system resources without causing undue "pain" to users, including "power users" and users who are subject to delays and input/output timing changes by non-ideal networks (which can cause otherwise normal user activity to appear to be unacceptable activity).

Rate Limiting and "Eviling"

[0046] The concepts of automatic rate limiting and modification of a user's "evil index" can be combined to provide a refined self-policing, automatic rate limiting system that can regulate an "eviled" user's ability to participate in forums without requiring total exclusion from the forum. Further, a user's current global or forum "evil index" can be used to modify the parameters used to effect automatic (i.e., "non-eviled" basis) rate limiting for the user. For example, a user's ability to interact in a forum can be rate limited by modifying the values for A , C , W , R , and/or D as a function of whether the user's "evil index" has been modified, or as a function of the current value of the user's "evil index". Conversely, the amount of modification of a user's "evil index" can be a function of the current value of A . Thus, conceptually, a user's permitted message rate R_p is a function of the user's "evil index" EI plus the user's rate of attempted message activity A : $R_p = f(EI) + g(A)$

[0047] Conceptually, each user's user-record thus may look like the following table:

Forum ID	"evil index"	decay rate	# times "eviled"	# times "eviled" others	time-of-entry	total-time of-access	A	C	W	R	D
Global value											
Forum 1 value											
Forum 2 value											

EP 1 124 173 A2

(continued)

Forum ID	"evil index"	decay rate	# times "eviled"	# times "eviled" others	time-of-entry	total-time of-access	A	C	W	R	D
etc.											

[0048] Separate "lifetime" and "session" records may be kept where useful. Of course, other or different information may be recorded for each user, and other ways of organizing such data may be used. Further, users may be assigned a "class" designation (such as "sysop" (system operator), corporate users, paying members, non-paying members, etc.) which can be given weight in making a determination of right to "evil" or absolute or relative immunity from being "eviled" in STEP 306 of FIG. 3.

[0049] In setting rate parameter values for a user, a global rate table can be used which is indexed by a user's "evil index"; that is, message rate parameters are a function of the user's global or forum "evil index". A conceptual example of one such table might be the following (ΔA represents an optional adjustment to be applied to the calculated value of A):

"evil index" value	ΔA	C	W	R	D
0-20					
21-40					
41-60					
61-80					
81-90					
91-95					
96-100					

[0050] The values for A, C, W, R, and D in this table can be defined globally or by forum. If by forum, the values can be defined in absolute terms or as relative offsets to a global table. Multiple tables of this form also can be defined, indexed by message type, so that misuse of certain types of messages are punished more than misuse of other message types.

[0051] Alternatively, a non-tabular implementation could be used instead of the global rate table described above. For example, a user's rate parameter values (ΔA , C, W, R, D) could be defined as a function describing a relation between these quantities.

Implementation

[0052] The methods and mechanisms described here are not limited to any particular hardware or software configuration, but rather they may find applicability in any computing or processing environment used in connection with online computer services.

[0053] The invention may be implemented in hardware or software, or a combination of both. However, preferably, the invention is implemented in computer programs executing on programmable computers each comprising at least one processor, at least one data storage system (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device. Program code is applied to input data to perform the functions described herein and generate output information. The output information is applied to one or more output devices, in known fashion.

[0054] Each program is preferably implemented in a high level procedural or object oriented programming language to communicate with a computer system. However, the programs can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted language.

[0055] Each such computer program is preferably stored on a storage media or device (e.g., ROM or magnetic diskette) readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage media or device is read by the computer to perform the procedures described herein. The inventive system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a specific and predefined manner to perform the functions described herein.

[0056] A number of embodiments of the present invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the scope of the invention as defined with reference to the following claims.

5

Claims

1. A method, performed on a computer system, of regulating a user's access to a computer-based service, the method comprising the steps of:
10 (a) computing (402) an average message rate for messages originated by the user's connection to the computer-based service;
 (b) comparing (404,408,412,420) the user's average message rate to a rate limit threshold (D,R,W,C);
15 (c) if the average message rate exceeds the rate limit threshold, then modifying (406,410,416,418) the user's ability to access the computer-based service.
2. The method of claim 1, further comprising the steps of:
 (a) comparing (404) the user's average message rate to a disconnect threshold;
20 (b) if the average message rate exceeds the disconnect threshold, then denying (406) the user access to the computer-based service.
3. The method of claim 1, further comprising the step of allowing (422) the user full access to the computer-based service if the average message rate falls below a clear threshold (C).
25
4. The method of claim 2, further comprising the step of allowing (422) the user full access to the computer-based service if the average message rate falls below a clear threshold (C).
5. The method of claim 4, further comprising the step of adjusting any of the average message rate, the rate limit threshold, the disconnect threshold, or the clear threshold in response to negative input about the user from at least one other user of the computer-based service.
30
6. The method of claim 4, wherein the step of computing (402) an average message rate for messages originated by the user's connection includes counting only messages of selected types.
35
7. The method of claim 4, wherein the step of computing (402) an average message rate for messages includes starting with an initial average message rate.
8. The method of claim 4, further comprising adjusting a value of the initial average message rate to vary the computer-based service's tolerance of initial messaging activity.
40
9. A method, performed on a computer system, of monitoring user activity in a computer-based service, the method comprising:
45 receiving input about a first user's activities from at least one other user of the computer-based service; and adjusting an index representative of the first user's activities based on the received input.
10. The method of claim 9, further comprising modifying the first user's ability to access the computer-based service based on the first user's index.
50
11. The method of claim 9, in which the input received from the at least one other user is in response to action taken by the first user.
12. The method of claim 9, further comprising limiting received input only to other users who have been affected by the first user's activities.
55
13. The method of claim 9, in which the index representative of the first user's activities corresponds to the appropriateness of the first user's behavior in using the computer-based service.

EP 1 124 173 A2

14. The method of claim 9, in which the input received from the at least one other user is sent anonymously.
15. The method of claim 14, in which adjustment of the first user's index comprises according reduced significance to anonymous input received from the at least one other user.
- 5 16. The method of claim 9, in which the input received from the at least one other user is sent automatically in response to action taken by the first user.
- 10 17. The method of claim 16, in which adjustment of the first user's index comprises according reduced significance to automatic input received from the at least one other user.
18. The method of claim 9, in which adjustment of the first user's index comprises according varied significance to received input based on a class to which the at least one other user belongs.
- 15 19. The method of claim 18, in which adjustment of the first user's index comprises according increased significance to input received from a user belonging to a system operator class.
- 20 20. The method of claim 18, in which adjustment of the first user's index comprises according reduced significance to input received from a user belonging to a non-paying member class.
21. The method of claim 18, in which adjustment of the first user's index comprises according varied significance to received input based on a seniority level of the at least one other user.
- 25 22. A computer program, residing on a computer-readable medium, for regulating access to a computer based service, comprising instructions for causing a computer to carry out the steps of any one of claims 1 to 21.
23. A method, performed on a computer system, of regulating a user's access to a computer-based service, the method comprising the steps of:
- 30 (a) receiving input about a first user from at least one other user of the computer-based service;
(b) evaluating the received input; and
(c) modifying the first user's ability to access the computer-based service based on a result of such evaluating.
- 35 24. The method of claim 23, in which the step of receiving input comprises gathering a vote from at least one other user of the computer-based service, the vote being cast in response to action taken by the first user.
25. The method of claim 24, in which the step of evaluating comprises applying predetermined access limiting criteria to the cast vote.
- 40 26. The method of claim 24, in which the step of evaluating further includes determining whether each other user is permitted to cast a vote against the first user.
27. The method of claim 26, in which the step of determining further includes evaluating whether a voting user has been accessing at least a portion of the computer-based service for a minimum amount of time.
- 45 28. The method of claim 26, in which the step of determining further includes evaluating whether a voting user has been accessing at least a portion of the computer-based service for a longer amount of time than the first user.
- 50 29. The method of claim 26, in which the step of determining further includes limiting the number of votes a user may cast.
30. The method of claim 24, in which the step of evaluating further includes giving extra weight to votes cast by another user against the first user if the other user has been accessing at least a portion of the computer-based service for a selected amount of time.
- 55 31. The method of claim 24, in which the step of modifying comprises denying the first user full access to the computer-based service based on the result of the evaluating.

EP 1 124 173 A2

32. The method of claim 31, further comprising the step of allowing the first user to regain full access to the computer-based service based on predetermined access resumption criteria.
- 5 33. The method of claim 23, further comprising the step of notifying selected other users that the first user's ability to access the computer-based service has been modified.
34. The method of claim 23, further comprising the step of providing an indicator of an amount of modification by all users accessing at least a portion of the computer-based service.
- 10 35. The method of claim 23, in which the computer-based service comprises an online public forum.
- 15 36. The method of claim 23, further comprising the step of permitting the first user to vote against another user based on a message type sent by such other user after the first user's ability to access the computer-based service has been modified.
- 20
- 25
- 30
- 35
- 40
- 45
- 50
- 55




FIG. 1
(PRIOR ART)




FIG. 2
(PRIOR ART)

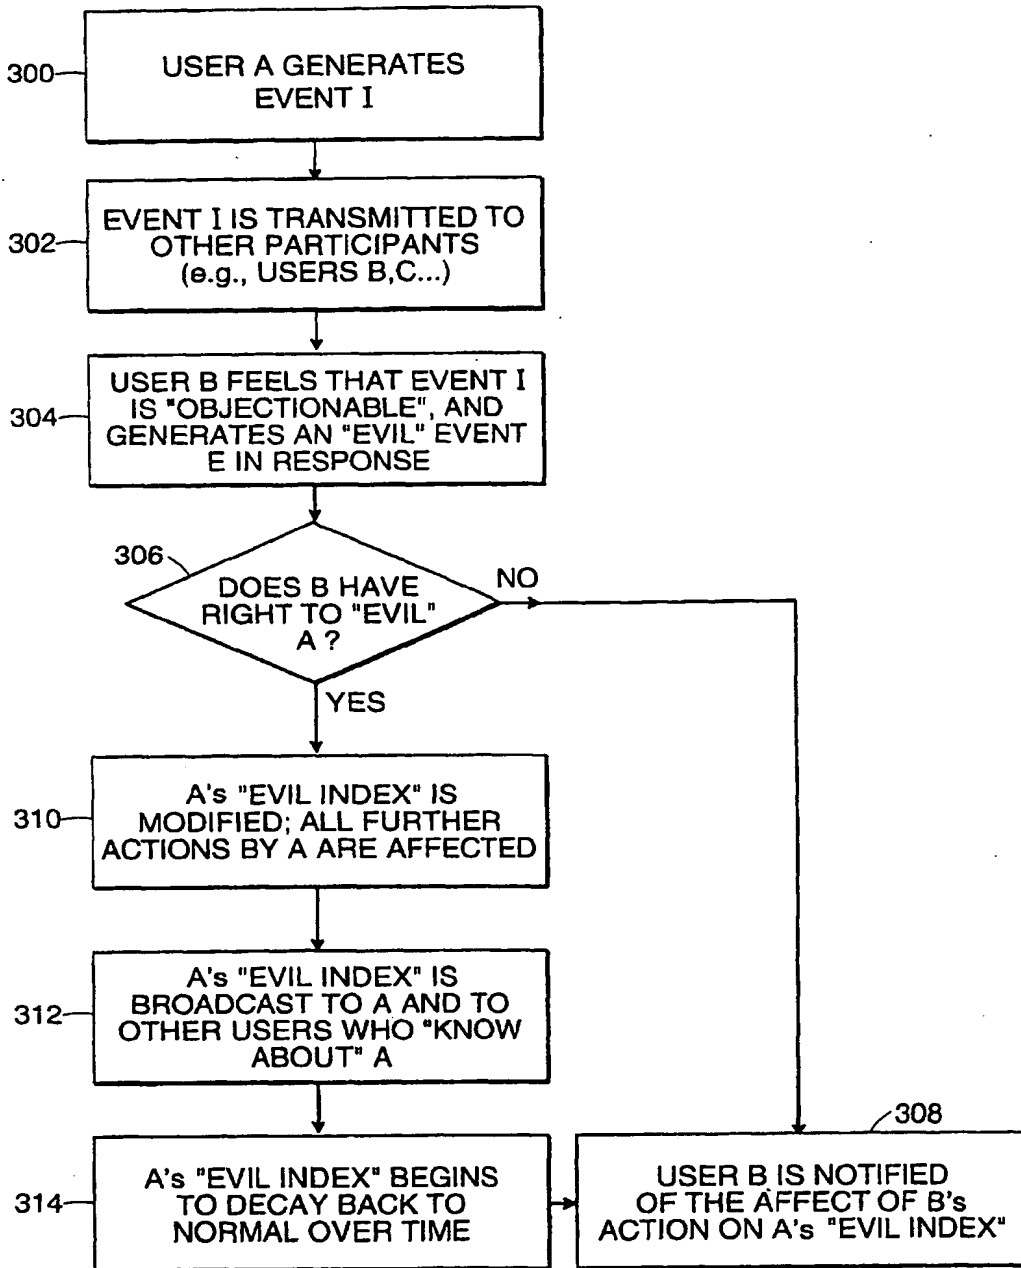


FIG. 3

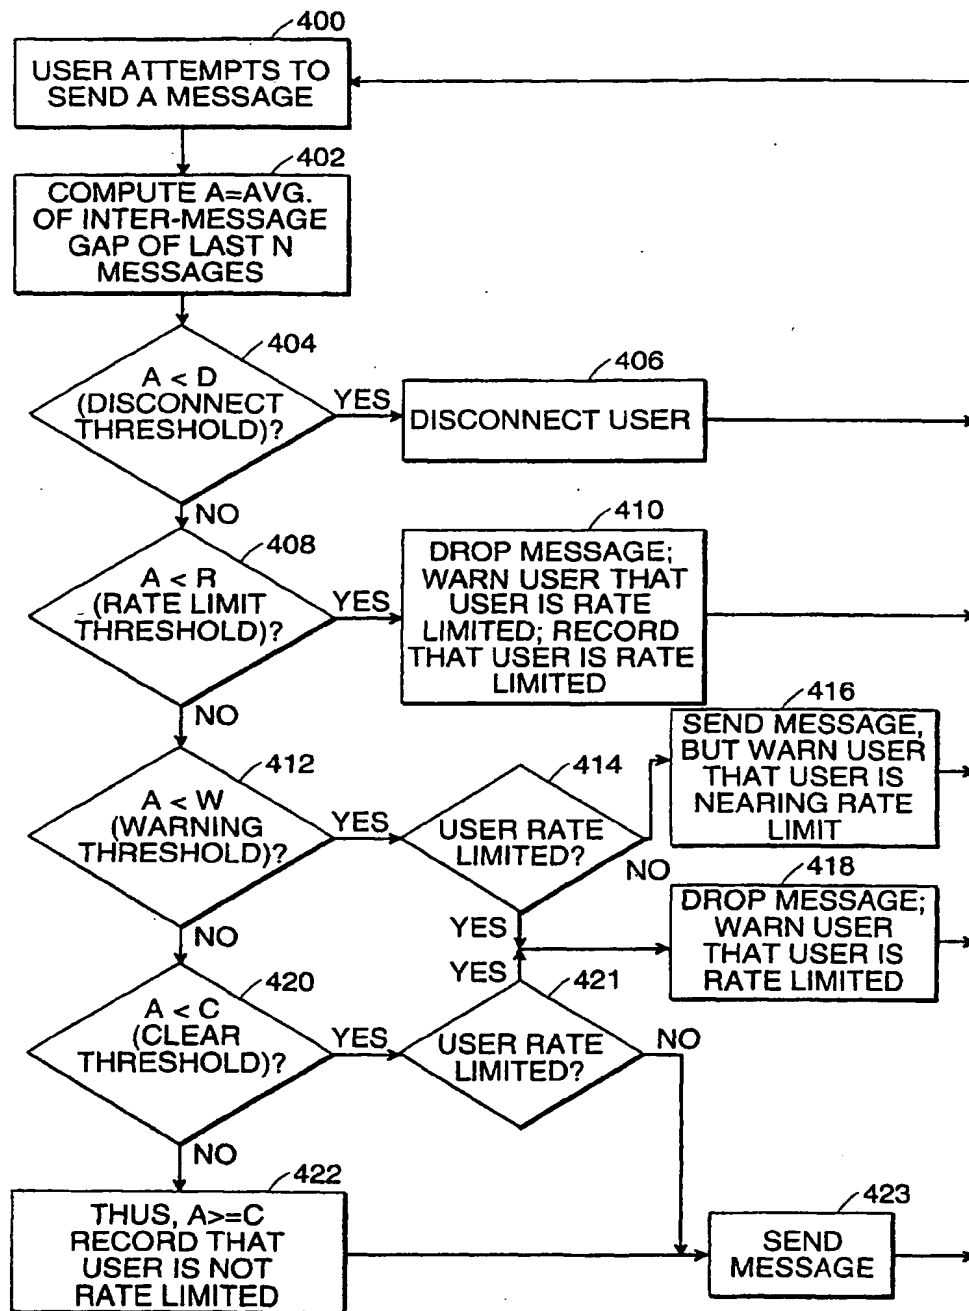


FIG. 4